

WE CLAIM:

1. A method for allocating a service on a network, said method comprising:
 - 5 collecting a set of performance data representative of a set of physical characteristics of the network;
identifying a plurality of node clusters in response to said collection of said set of performance data; and
correlating said plurality of node clusters and a set of at least one
10 performance rule for said plurality of node clusters as related to the service.
2. The method of claim 1, further comprising:
providing a map as a result of said correlation, said map including a
first cluster of said plurality of clusters for supporting the service on the network.
15
3. The method of claim 2, further comprising:
allocating the service to said first cluster.
4. The method of claim 1, further comprising:
20 providing a listing as a result of said correlation, said listing including at least one server within a first cluster of said plurality of clusters for supporting the service on the network.
5. The method of claim 4, further comprising:
25 allocating the service to a first server of said at least one server.

6. The method of claim 1, further comprising:
probing the network for a round trip time.
- 5 7. The method of claim 1, further comprising:
probing the network for a hop count.
8. The method of claim 1, further comprising:
probing the network for a bottleneck link speed.
- 10 9. A distributed computing system, comprising:
a plurality of interconnected nodes; and
a server operable to allocate a service for said plurality of
interconnected nodes, said server including
- 15 a probe operable to provide a set of performance data as related to
a set of physical characteristics of said plurality of interconnected nodes,
a module operable to identify a plurality of node clusters within said
network in response to said set of performance data; and
- 20 an engine operable to utilize at least one performance rule for said
plurality of node clusters as related to said service to identify a first node cluster
of said plurality of node clusters for supporting said service for said plurality of
interconnected nodes.
10. The system of claim 9, wherein
- 25 a round trip time of said plurality of interconnected nodes is a first
performance data of said set of performance data.

11. The system of claim 9, wherein
a hop count of said plurality of interconnected nodes is a first
performance data of said set of performance data.

5

12. The system of claim 9, wherein
a bottleneck link speed the plurality of interconnected nodes is a
first performance data of said set of performance data.

10

13. The system of claim 9, wherein
said engine is further operable to correlate said plurality of node
clusters and said at least one performance rule to thereby identify said first node
cluster.

15

14. The system of claim 9, wherein
said module is a neural network.

15. A computer program product in a computer readable medium for
allocating a service on a network, comprising:

20

a means for collecting a set of performance data relating to a set of
physical characteristics of a network;

a means for identifying a plurality of node clusters in response to
said set of performance data; and

a means for correlating said plurality of node clusters and a set of
at least one performance rule for said plurality of node clusters as related to said
service.

25

16. A server including a memory and a processor for allocating a service on a network having a plurality of interconnected nodes, comprising:

- 5 a probe operable to provide at least one performance data as related to a set of physical characteristics of the plurality of interconnected nodes,
- a module operable to provide a plurality of node clusters of the network in response to said set of performance data; and
- 10 an engine operable to utilize at least one performance rule for said plurality of node clusters as related to the service to identify a first node cluster of said plurality of node clusters for supporting the service for the plurality of interconnected nodes.

17. The server of claim 16, wherein
- 15 a round trip time of the plurality of interconnected nodes is a first performance data of said set of performance data.

18. The server of claim 16, wherein
- 20 a hop count of the plurality of interconnected nodes is a first performance data of said set of performance data.

19. The server of claim 16, wherein
- a bottleneck link speed of the plurality of interconnected nodes is a first performance data of said set of performance data.

25

20. The server of claim 16, wherein
- said module is a neural network.

21. The server of claim 16, wherein
- said engine is further operable to correlate said plurality of node clusters and said at least one performance rule to thereby identify said first node
- 5 cluster.